PART VII

APPLICATIONS OF PROBABILITY: PHILOSOPHY

CHAPTER 29

PROBABILITY IN EPISTEMOLOGY

MATTHEW KOTZEN

29.1 Introduction

Increasingly in recent years, formal approaches to epistemology have become prevalent, especially ones that involve some sort of appeal to probabilities. Typically, advocates of these approaches do not claim that formal approaches could or should entirely replace traditional epistemological approaches; rather, they usually claim only that formal tools can be used to supplement traditional techniques in order to precisify various epistemological problems, theses, and arguments. In this article, I survey some of the areas of epistemology in which appeals to probability have been most influential.

29.2 Full and Partial Beliefs

The "traditional" concept of a belief is a binary concept; for all *A*, one either believes *A* or fails to believe *A*. Much of traditional epistemology, then, focuses on the question of what it takes for a belief that *A* to be *justified*, or to be an item of *knowledge*, etc. Another way of conceiving of belief is as a *graded* notion; on this conception, one's doxastic attitude toward *A* can be represented by a real number in the [0,1] interval—referred to as one's "credence" in *A*—where higher credences correspond to "more" belief in *A*. There are different stories that we might tell about how these conceptions of belief are related; I will discuss four such views.

On one view, there are psychological states that correspond to binary beliefs, and there are different psychological states that correspond to credences; thus, the two conceptions are not competitors for the correct account of the metaphysics of doxastic attitudes, but are rather accounts of two different sorts of doxastic attitudes. But there is clearly *some* metaphysical connection between these attitudes; it would be hard to make sense of an agent who claimed to have credence 0 in a proposition *A* and yet claimed to believe that *A*. Moreover, some philosophers have worried that this view simply isn't psychologically realistic, since "[t]here is no evidence to believe that the mind contains two representational systems, one to represent things as being probable or improbable and the other to represent things as being true or false."[1]

On a second view of the relation between binary beliefs and credences, credences are to be analyzed as full beliefs in explicit claims about objective probabilities; thus, for example, a credence of .6 in the proposition that it is going to rain tomorrow is to be analyzed as a full belief in the proposition that the objective probability of rain tomorrow is 60%.[2] But it's not clear that there is a coherent notion of objective probability that can support this analysis in all cases; for example, I might have a credence of .7 that Tom was at the party yesterday, but it's at least somewhat natural to think that the objective probability that Tom was at the party yesterday is either 1 or 0, depending on whether he was there or not. Similarly, I might have a credence of .7 in the proposition that a particle will decay in the next hour because I'm 50% sure that the particle has an objective probability of 80% of decaying, and 50% sure that the particle has an objective probability of 60% of decaying; in such a case, though my credence that the particle will decay is .7, I definitely do *not* believe that the objective probability that the particle will decay is 70%. Again, it's not clear that the view under consideration can be extended to this sort of case. Someone might respond to both of these cases by pointing out that there is a coherent *evidential* notion of objective probability on which the objective probability of Tom's being at the party yesterday is .7, and on which the objective probability of the particle decaying is .7.[3] But this doesn't seem to lead to a plausible theory of the relation between binary beliefs and credences either; my credence that Tom was at the party yesterday still might be .7, even though I don't have *any* beliefs at all about evidential probabilities.

On a third view, binary belief is to be analyzed in terms of credence; on the most natural way of developing this view, to have a binary belief that A is to have a credence in A that is above a certain threshold, where that threshold is determined at least in part by either the context of the believer or the context of the belief-ascriber. Kaplan 1996 refers to this view as the "threshold view." Two significant problems arise for this view, both of which are discussed in Stalnaker 1984. First, it is hard to see which features of either the believer's or the belief-ascriber's context could make it the case that the threshold value for belief is, say, .9786 rather than .9785. Some philosophers (such as Foley 1993 chapter 4 and Hunter 1996) have argued that the threshold must therefore be vague, but it is not completely clear whether this helps.[4] Second, if an agent regards A and B to be independent, then it is possible for her to have a rational credence in A above a given threshold, and also to have a rational credence in B above that threshold, and yet to have a rational credence in $A \wedge B$ that is below the threshold. If binary belief is just credence above the relevant threshold, then this situation corresponds to one in which the agent believes A, and also believes B, but fails to believe $A \wedge B$. This is somewhat odd, since we usually think of an agent who believes A and B separately but who fails to believe $A \wedge B$ as exhibiting some rational failing. This issue has received much attention in discussions of the closely-related Preface Paradox—see, e.g., Christensen 2004—and the Lottery Paradox—see, e.g., Kyburg 1961.

[1] Weatherson 2005, p. 420.
[2] See, e.g., Frankish 2009, p. 77.
[3] This notion of "evidential probability" is discussed in Section 6.
[4] See, e.g., Weatherson 2005, pp. 420–1.

On a fourth view, credences are understood to *replace* binary beliefs. On this view, there just are no psychological states corresponding to binary beliefs; rather, all that we have are credences in various propositions—some higher and some lower—but none of those credences stand out for special classification as our "beliefs."[5] Of course, it's often convenient to talk in terms of full belief, as when we say "I believe it's going to rain tomorrow," but strictly speaking, there are no psychological states that correspond to the binary notion. This view naturally raises the question of how to reconceive traditional epistemological concepts such as *knowledge* in the context of credences. Moss [forthcoming] develops an answer to this question.

29.3 Synchronic Constraints

So far, the only constraint on credence functions that we have seen is the constraint that they are real numbers in the [0,1] interval. There is a variety of different further constraints that have been argued for. Some of these constraints are *synchronic*; they are constraints that are alleged to apply to agents *at a particular time*. Other constraints are *diachronic*; they apply to agents' *revisions* of their credences from one time to another. This section holds a discussion several proposed synchronic constraints on credences, and Section 4 discussion of several proposed diachronic constraints.

Even within the synchronic constraints on credences, we can distinguish constraints on our *unconditional* credences from constraints on our *conditional* credences. Unconditional credences are credences that we assign to individual propositions. Conditional credences, by contrast, are often expressed as credences in A, *given* B. Traditionally, a conditional credence $p(A|B)$ was understood to be defined as $\frac{p(A\wedge B)}{p(B)}$. More recently, Alan Hájek has argued that we should instead understand conditional credences as primitive; among other things, this understanding has the advantage of allowing conditional credences to be defined when $p(B) = 0$ (which the traditional definition clearly does not).[6] But the issue of how conditional credences are best to be understood is in principle separable from the question of how they are rationally constrained (even if these questions turn out to be related in various ways). Sections 3.1 and 3.2 develop some constraints on unconditional credences, and Sections 3.3 and 3.4 develop some constraints on conditional credences.

29.3.1 Probabilism

One of the most widely discussed synchronic constraints on unconditional credence is the constraint that credences obey the standard Kolmogorov probability axioms.[7] This constraint can be developed as follows: we start with a set of mutually exclusive and jointly exhaustive possibilities, which is designated Ω. Then, we define a field F of subsets of Ω; call

[5] Jeffrey 1970 (pp. 171–72) expresses sympathy for this view: "I am inclined to think Ramsey sucked the marrow out of the ordinary notion [of belief], and used it to nourish a more adequate view."

[6] Hájek 2003.

[7] Kolmogorov 1956.

the members of F "propositions." Finally, we define a probability function p from F to the [0,1] interval of the real numbers that obeys the following three axioms for all propositions A and B in F:

1. $p(A) \geq 0$.
2. $p(\Omega) = 1$.
3. If $A \cap B = \emptyset$, then $p(A \cup B) = p(A) + p(B)$.

The distinctive claim of Probabilism is that it is a rational constraint on credences that they be probability functions, in the sense defined above. So, according to Probabilism, a rational agent must (1) assign a non-negative credence to every proposition, (2) assign a credence of 1 to the set of all possibilities, and (3) where A and B are mutually exclusive, assign a credence to $A \vee B$ that is the sum of the credences she assigns to A and to B separately. A credence function obeying these constraints is called "coherent," and one not obeying these constraints is called "incoherent."[8]

A variety of different arguments have been offered for Probabilism. First, Synchronic Dutch Book Arguments proceed from the premise that agents with incoherent credences are committed in some way (in virtue of the Dutch Book Theorem) to a series of bets that jointly guarantee them a loss; but since (the Argument continues) being so committed is irrational, it must be a rational constraint on credences that they be probabilities.[9] Early Synchronic Dutch Book Arguments assumed that the nature of this "commitment" was one of identity; according to these arguments, to have a credence in A of .5 *just is* to be disposed to accept a bet at 1 : 1 odds that wins if and only if A is true. On this understanding, the Dutch Book Theorem shows that agents with incoherent credences are disposed to accept each in a series of bets such that the bets jointly guarantee that the agent will lose money. This version of the Synchronic Dutch Book Argument faced the objection that a rational agent might well have a credence in A of .5, say, and yet not be disposed to bet on A at 1 : 1 odds owing to a distaste for gambling, or owing to risk-aversion, etc. Thus, more recently, there have been a variety of "depragmatized" Dutch Book Arguments which claim only that an agent with a credence in A of .5 is *committed to the fairness of* a bet at 1 : 1 odds, regardless of whether he is actually disposed to accept such a bet.[10] According to these arguments, it is a rational failing to be committed to the fairness of a series of bets that jointly guarantee a loss; hence, since incoherent credences commit an agent in precisely this way, rational agents are constrained to have coherent credences.[11]

A second type of argument for Probabilism appeals to a "representation theorem."[12] Such theorems establish that, if an agent's preferences satisfy various plausible constraints such as transitivity and connectedness, then that agent is representable as having a utility function $U(x)$ and a coherent credence function $C(x)$ such that the agent prefers A to

[8] For a general discussion of this topic, see Lyle Zynda's article "Subjectivism" in this volume (Zynda 2014).

[9] The original Synchronic Dutch Book Argument is usually attributed to Ramsey 1926 and de Finetti 1937.

[10] See Armendt 1993, Christensen 1996, and Skyrms 1984.

[11] See Easwaran 2011 for an excellent overview of Dutch Book Arguments.

[12] See Savage 1954, Jeffrey 1965, Joyce 1999, and Maher 1993.

B just in case her "expected utility" of *A*, calculated via $U(x)$ and $C(x)$, is higher than her expected utility *B*. Arguments for Probabilism based on representation theorems have faced a variety of objections. Many have focused on the worry that just because a rational agent's preferences are *representable* as being derived from a utility function $U(x)$ and a coherent credence function $C(x)$, it doesn't obviously follow that he *actually has* $C(x)$ as his credence function; in other words, even if the representation theorem argument establishes that all rational agents are representable as having coherent credence functions, that conclusion falls short of establishing the Probabilist claim that all rational agents *actually do have* coherent credence functions.[13] In particular, Zynda argues that alternative representations of an agent who satisfies the intuitive constraints—ones which *don't* appeal to a coherent credence function—are possible, so the significance of the fact that such agents can *also* be represented as having coherent credence functions isn't clear. Other objections have included criticism of the rational constraints on preferences (such as transitivity and connectedness) that give rise to the representation theorem.[14]

A third argument for Probabilism appeals to the notion of *calibration*.[15] To have credences that are perfectly calibrated is to have credences that perfectly match relative frequencies; proportion *x* of the cases in which you assign a credence of *x* to a proposition are cases in which that proposition is true, for all *x*. (So, 20% of the cases in which you assign a credence of .2 to a proposition are cases in which the relevant proposition is true, 30% of the cases in which you assign a credence of .3 to a proposition are cases in which the relevant proposition is true, etc.) However, even if you are not *perfectly* calibrated, it still makes sense to talk about being *more* or *less* calibrated; for instance, other things equal, if *A* is true in 21% of the cases in which I assign a credence of .2 to *A*, my credence function is better calibrated than it would be if *A* were true in 22% of the cases in which I assign a credence of .2 to *A*. The argument for Probabilism, then, shows that if your credence function is incoherent, then there is some coherent credence function that is better calibrated than yours under any logically consistent assignment of truth-values to propositions; from this fact, it follows, allegedly that you're rationally constrained to have a coherent credence function.[16] However, some philosophers have expressed doubt whether perfect calibration is really a rational ideal.[17] Moreover, one might think that what really matters is how well-calibrated your credence function is in the *actual* world, not how well-calibrated it is in *all possible worlds*, and coherence does not guarantee maximal calibration in the actual world.[18]

Finally, a fourth argument for Probabilism, due to Joyce 1998, alleges that for any incoherent credence function, there corresponds a coherent credence function that is strictly more accurate under every logically consistent assignment of truth-values to propositions. This argument is similar in some ways to the calibration argument above, but whereas the calibration argument appeals to relative frequencies, Joyce's argument appeals to the *accuracy* of credences. Non-extreme credences, of course, can't be accurate in the same way that binary beliefs can be (since a binary belief is accurate iff it is true, and there's no

[13] See, e.g., Hájek 2008.
[14] See, e.g., Schick 2003.
[15] See van Fraassen 1983 and Shimony 1988 for examples of this argument strategy.
[16] See Joyce 2004 for an explanation of this style of argument.
[17] See Seidenfeld 1985, Joyce 1998, and Hájek unpublished A.
[18] Hájek 2008.

obvious sense in which a non-extreme credence can be true); however, as in the calibration argument, we can say that when *A* is true, a credence in *A* is more accurate the closer it is to 1. Joyce's argument then proceeds to show that incoherent credence functions are always "less accurate than they could be," since whenever an agent's credence function is incoherent, he has available to him some coherent credence function which is more accurate than his credence function, no matter how the world turns out to be. Important objections to Joyce's argument are raised in Bronfman unpublished, Hájek 2008, and Maher 2002; some of these objections parallel the objection discussed above to the calibration argument by arguing that what really matters is how accurate your credence function is in the *actual* world, rather than *all possible worlds*. Joyce responds to these and other objections in Joyce 2009.

29.3.2 Regularity

Another proposed rational constraint on unconditional credence is the constraint that, if an agent regards a proposition *A* to be possible, then she should assign positive credence to *A*. The nature of the possibility at issue here is controversial; some authors have logical possibility in mind, whereas others seem to appeal to metaphysical possibility.[19]

But however we understand it, this constraint isn't entailed by the probability axioms; all that axiom 1 from Section 3.1 requires is that the agent assign *non-negative* credence to each proposition. And it's at least very plausible that an agent is permitted (perhaps even *required*) to assign a credence of 0 to propositions that she regards to be *impossible*: the proposition that she doesn't exist, for instance, or the proposition that 1+1=17. The constraint under consideration here entails that propositions that an agent regards to be impossible are the *only* propositions to which she is rationally permitted to assign a credence of 0. That constraint is called Regularity.[20]

On the one hand, Regularity can seem plausible. It's natural to think that your credence function should distinguish between those propositions that you regard as possible and those that you regard as impossible; however, if there are propositions of each type to which you assign credence 0, then it looks as if your credence function isn't able to make this distinction in all cases. Similarly, if I believe that *B* is true in only a proper subset of the worlds in which *A* is true, then it is natural to think I ought to assign a higher credence to *A* than to *B*; after all, I think that there are ways for *A* to be true while *B* is false, but no ways for *B* to be true while *A* is false. But this natural thought looks to entail Regularity. If, for any *A* and *B*, my believing that there are possibilities compatible with *A* but not with *B* (and not vice versa) rationally requires my assigning a higher credence to *A* than to *B*, then that must be because I am required to assign positive credence to the (non-empty) set of possibilities in which *A* is true but *B* is false, but I am required to do that only if I am required to assign positive credence to every proposition that I regard as possible.

[19] The versions of Regularity discussed in Shimony 1970 and Skyrms 1980 appeal to the former understanding of possibility, whereas Lewis 1980 (at least arguably) appeals to the latter. For an overview of the options, see Hájek 2012 and Hájek unpublished B.

[20] Versions of the Regularity constraint have been proposed by Kemeny 1955, Shimony 1955, Shimony 1970, Jeffreys 1961, Edwards et al. 1963, Carnap 1963a, Stalnaker 1970, Lewis 1980, Skyrms 1980, Appiah 1985, Jackson 1987, and Jeffrey 1992.

Regularity runs into trouble, however, in certain cases with uncountably many possible outcomes. Suppose that I am about to throw an infinitely sharp dart at a dartboard, and suppose for simplicity that the dart is guaranteed to hit a sport somewhere on the dartboard. Now, consider the uncountable set of all of the points on the dartboard; for each such point, I regard the proposition that the dart lands on precisely that point to be possible. But there is no way for a coherent credence function to assign a (real-valued) positive credence to each of these propositions; indeed, as Hájek 2003 shows, any probability function defined on an uncountable algebra assigns probability 0 to uncountably many propositions. A similar problem arises for unmeasurable sets of points. Consider some unmeasurable set of points on the dartboard *S*, and consider the proposition *B* that the dart lands on one of the points in *S*. As before, I regard *B* to be possible; after all, the dart could certainly hit any individual point in *S*, in which case *B* would be true. But again, there is no way for a coherent credence function that respects the symmetries of the situation to assign a (real-valued) positive credence to *B*.[21] In response to these complications, some have proposed a relaxation of the requirement that credence functions be real-valued, suggesting instead that we should allow hyperreal-valued credence functions.[22] Williamson 2007 and Hájek unpublished B raise various objections to this proposal.

29.3.3 Reflection

Another alleged synchronic constraint on rationality is van Fraassen's Reflection Principle."[23] Here is van Fraassen's statement of the "General Reflection Principle": "My current opinion about event E must lie in the range spanned by the possible opinions I may come to have about E at later time *t*, as far as my present opinion is concerned."[24] This constraint is naturally understood as a constraint on an agent's conditional credences: her current credence in *A*, conditional on the proposition that her future credence in *A* will be in the [m,n] range, should itself be somewhere in the [m,n] range.

This principle is motivated by the thought that you are rationally obligated to treat your future self as a better authority with regard to the subject matter of your beliefs than you are. So, conditional on the information that your future self is going to have a credence in *A* that is in some range, it can seem odd to assign a credence to *A* that is outside that range rather than to "respect" your future self's epistemic authority by assigning a conditional probability that conforms to hers. As a result, the thought goes, even if you aren't actually certain of what your future credences will be, you are still constrained to have *conditional* credences in each proposition that obey General Reflection.

Of course, General Reflection is much less plausible in cases where an agent rationally believes that at some time between now and the time at which her future self has a credence in *A* that is in the [m,n] range, she is going to lose some of her current evidence, or else

[21] See Hájek unpublished B for a discussion of complications arising from unmeasurable sets and sets of measure 0.

[22] See, e.g., Lewis 1980, who cites Bernstein and Wattenberg 1969 for their technical construction involving infinitesimals.

[23] See van Fraassen 1984 and van Fraassen 1995.

[24] van Fraassen 1995, p. 16.

become irrational or suffer some other sort of epistemic defect; for instance, I might have a rational conditional credence in *A*, conditional on my getting kidnapped tonight and brainwashed into having a credence in *A* of .8 by tomorrow, of .5.[25] But in cases where an agent rationally believes that she won't have lost any evidence that she currently has or suffer any other sort of epistemic deficit by *t*, it is more plausible that she is rationally constrained to obey General Reflection. This constraint clearly isn't entailed by the agent having a probabilistic credence function, so van Fraassen thinks that Reflection constitutes a substantial additional synchronic constraint on rationality.

Various objections have been raised against Reflection. Elga 2000 argues for the "Thirder" answer to the Sleeping Beauty Problem (discussed in Section 4.2 below), and claims that that answer is incompatible with Reflection. Arntzenius 2003 discusses several cases where it seems as though a rational agent violates Reflection even though no *actual* epistemic deficit befalls the agent; in Artnzenius's cases, it seems to be enough that the agent's future self will have some positive credence that the relevant deficit has been suffered (and that the agent's previous self knows this), even if no deficit is ever actually suffered by the agent.

29.3.4 The Principal Principle

In Lewis 1980, David Lewis defended a synchronic rational constraint on conditional credences that, he argues, is imposed by information about objective chances.[26] The guiding intuition behind this constraint is that a rational agent's conditional credence at *t* that *A* is true, conditional on the information that *A*'s objective chance at *t* is *x*, should itself be equal to *x*; for instance, a rational agent's conditional credence at 1:00 that the coin will land heads at 2:00, conditional on the information that the chance at 1:00 that the coin will land heads at 2:00 is .423, should itself be .423.

However, Lewis claims, a rational agent's conditional credence at *t* that *A* is true, conditional on the information that *A*'s objective chance at *t* is *x and* that *E* is true, might not be *x*; in cases where *E* contains "inadmissible" information, the agent's conditional credence might take a value other than *x*. For Lewis, admissible information is "the sort of information whose impact on credence comes entirely by way of credence about the chances of those outcomes"[27]; a prime example of inadmissible information is information about the future. For instance, suppose that *E* is the information that a known-to-be-reliable crystal ball predicts, at 1:00, that the coin will land heads at 2:00. Then, the Principal Principle doesn't entail that a rational agent's conditional credence at 1:00 that the coin will land heads at 2:00, conditional on the information that the objective chance at 1:00 that the coin will land heads at 2:00 is .5 *and* that *E* is true, is .5; in this case, the rational agent's conditional credence might be 1, for instance.

[25] See Christensen 1991 and Talbott 1991 for discussion of cases like these.

[26] Lewis 1980 contains the canonical formulation of this type of constraint, though Lewis there acknowledges his debt to Mellor 1971, which (in Lewis's words) "presents a view very close to" Lewis's own view.

[27] Lewis 1980, p. 92.

Following Lewis 1980: Let p be any "reasonable initial credence function."[28] Let x be any real number in the [0,1] interval. Let X be the proposition that the objective chance at t of A's holding is x. And let E be any proposition that is both compatible with X and admissible at t. Then, the Principal Principle says that p will satisfy the constraint that $p(A|X \wedge E) = x$. In other words: as long as an agent's initial credence function is reasonable, her initial conditional credence in A, conditional on the conjunction of a proposition specifying the objective chance of A at t with any (compatible) admissible proposition, is just the value that was specified to be the objective chance of A at t.

Later, Lewis worried that the Principal Principle conflicted with his own Humean "best system" theory of laws and chances, owing to the so-called "problem of undermining futures."[29] The source of this problem is that, if a Humean theory of chance—according to which present chances supervene on the whole of history, including the future as well as the present and the past—is true, then at any time t, there are non-actual futures with nonzero chances which are such that, were they to transpire, some of the chances at t would have different values from the ones they actually have. For instance, suppose that a coin is flipped a large but finite number of times in the entire history of the universe, and consider some time t, by when only some small fraction of the flips have taken place. Suppose that the coin actually lands heads (approximately) half the time in the entire history of the universe, so that the best theory of chance entails that the chance of any particular coin flip landing heads is .5. Still, according to that theory, the "undermining future" F in which the coin land heads on each flip after t has some positive chance; and if such a future were to obtain, then the best theory of chance would entail that the chance at t of any particular coin flip after t landing heads is greater than .5 (and hence different from its actual value of .5).

So far, Lewis thinks that this consequence is merely "peculiar," but the conflict with the Principal Principle arises when we consider the chance at t of F itself. The actual chance at t of any particular flip after t landing heads is, we're supposing, .5. Suppose that there are n flips which occur after t. Then, the actual chance at t of F is $.5^n$. But if F were to occur, then the chance at t of any particular flip after t landing heads would be greater than .5, and hence the chance of t of all n flips after t landing heads would be m^n, for some $m > .5$. Let X be the proposition specifying the actual chance of F at t (i.e., $.5^n$), and let E be any admissible proposition compatible with X. By the Principal Principle, a rational agent's $p(F|X \wedge E)$=$.5^n$, which is positive (since n is finite). But, since F entails a different chance at t of F than X entails, F is inconsistent with X, and hence a rational agent's $p(F|X \wedge E) = 0$. Contradiction.

One possible response would be to say that information about objective chances (such as X in the previous paragraph) is itself information about the future, and hence inadmissible.

[28] Lewis requires that a reasonable initial credence function at least obey the axioms of the probability calculus (see Section 3.1 of this chapter) and that it be regular (see Section 3.2 of this chapter). He also says that a reasonable initial credence function is such that "if you always learned from experience by conditionalizing on your total evidence, then no matter what course of experience you might undergo your beliefs would be reasonable for one who had undergone that course of experience." (Lewis 1980, p. 268.) See Section 4.1 of this chapter for a discussion of conditionalization.

[29] Lewis develops this worry in Lewis 1986. Halpin 1994, Hall 1994, and Thau 1994 all advocate versions of the strategy of rejecting the Principal Principle in response to the problem. See Wolfgang Schwarz's article "Best Systems Approaches to Cnance (20)" in this volume (Schwarz 2014), for further discussion.

But, as Lewis points out,[30] this would render the Principal Principle useless; the whole point of the Principal Principle is to articulate the rational constraint imposed by information about current objective chances, and if Lewis's Humean theory of chances is correct, then current chances supervene in part on future events.

In Lewis 1994, Lewis qualifies the Principal Principle and endorses a modified "New Principle." Recall that the original Principal Principle says that, where X is the proposition that the objective chance at t of A's holding is x, a rational agent's $p(A|X \wedge E) = x$. The New Principle modifies the content of X; on the modified understanding, X is the proposition that the objective chance at t of A's holding, *conditional on the actual true theory of chance*, is x. We therefore ignore undermining futures which would contradict the actual true theory of chance.[31] However, some authors have argued that Lewis didn't actually need to abandon the original Principal Principle in favor of the New Principle.[32]

29.4 Diachronic Constraints

29.4.1 Conditionalization

Perhaps the most widely discussed diachronic constraint on credences is the Rule of Conditionalization. According to this rule, when you learn E and nothing more, you should update your credence in any proposition A from your old unconditional credence $p(A)$ to your old conditional credence $p(A|E)$. It is somewhat tempting to read this constraint as a trivial one, since it is tempting to understand $p(A|E)$ as "the credence that it's rational to assign to A, given that you've learned that E is true." But this is a mistake; $p(A|E)$ is an entirely synchronic feature of your old credence function, and doesn't by itself entail anything about how you should *update* your credence function when you acquire new evidence such as E. On the "traditional" view of conditional credences discussed in Section 3, $p(A|E)$ is just an abbreviation for $\frac{p(A \wedge E)}{p(E)}$; and it is clearly a substantive constraint on rationality that, when a rational agent learns that E, her new credence in A should be the ratio of her old credence in $A \wedge E$ to her old credence in E. And even if Hájek is right that conditional credences should be regarded as primitive rather than defined in the manner above, it is *still* a substantive constraint on rationality that, when you learn E, your *new* credence in A should be equal to your *old* primitive conditional credence $p(A|E)$.

One prominent argument that is offered in support of the Rule of Conditionalization is the Diachronic Dutch Book Argument, which—though similar in strategy—is importantly different from the Synchronic Dutch Book Argument discussed in Section 3.1. The first Diachronic Dutch Book Argument appeared in Teller 1973, though Teller attributed the argument to David Lewis.[33] Whereas the Synchronic Dutch Book Argument assumes that an agent who has a $p(A) = x$ is committed to regarding as fair a bet on A at $x : 1 - x$ odds,

[30] Lewis 1994, pp. 485-6.
[31] See Loewer 2004 and Roberts 2001 for clear expositions of the New Principle.
[32] See, e.g., Vranas 1998 and Roberts 2001.
[33] Lewis later developed a version of the Diachronic Dutch Book Argument in Lewis 1999.

the Diachronic Dutch Book Argument makes the additional assumption that an agent with $p(A|B) = x$ is committed to regarding as fair a *conditional* bet which is "called off" if B is false, and which becomes an ordinary bet on A at $x : 1 - x$ odds if B is true. Then, the argument proceeds by showing that an agent who violates the Rule of Conditionalization is committed to regarding as fair a series of bets, some made before learning B and some made after learning B, which jointly guarantee a net loss for the agent over time. In general, philosophers have found Diachronic Dutch Book Arguments to be less persuasive than their Synchronic cousins.[34]

Another important argument for the Rule of Conditionalization is developed in Greaves and Wallace 2006. The strategy of this argument is to define a notion of the "epistemic utility" of a credence according to which, on the assumption that A is true, a credence of x in A has higher utility than a credence of y in A just in case $x > y$. (On the assumption that A is false, a credence of x in A has more epistemic utility than a credence of y in A just in case $y > x$.) Then, Greaves and Wallace define an agent's *expected* epistemic utility, which depends on both the agent's credence distribution and the epistemic utilities of those credence functions. Greaves and Wallace then prove that under certain conditions, the Rule of Conditionalization is the unique update rule that maximizes expected epistemic utility.

There are some approaches to epistemic rationality that reject the need for any update rule at all. For example, in Williamson 2010, Jon Williamson defends a version of "Objective Bayesianism" that features three synchronic constraints on a rational agent's credences. But instead of endorsing any update rule that constrains an agent to *change* his credences in any particular way, Williamson instead argues that the three synchronic constraints should just be applied to each new evidential situation that the agent finds herself in. Thus, different evidential situations constrain an agent to have different credences, but not because new evidence mandates a *change* in credences *as such*; rather, new evidence creates a new evidential situation in which the three synchronic norms generate new synchronic constraints.

29.4.2 Updating on *De Se* Information

On the traditional understanding, credences are interpreted as being assigned to *uncentered* propositions, or propositions about what the world is like that reflect a "third-personal" perspective; we then address both synchronic questions about how an agent should assign his credences to various uncentered propositions, and also diachronic questions about how he should update his credences in those uncentered propositions as he collects new evidence. But in recent years, there has been considerable interest in the question of how an agent should apportion and update the credences he assigns to *centered* propositions, or propositions which reflect a "first-personal" perspective by being about the location or identity of the believer of the proposition.[35] It's plausible that to know which uncentered propositions are true isn't always to know everything about your situation, since you might

[34] See Christensen 1991 for one line of resistance to Diachronic Dutch Book Arguments.

[35] The seminal articles on centered possible worlds are Quine 1969, Perry 1979, Lewis 1979, and Stalnaker 1981. The rebirth of interest in the topic began with Elga 2000 and Lewis 2001. For a general discussion, see Mike Titelbaum's article "Self-Locating Credences" in this volume (Titelbaum 2014).

not know *who* or *where* or *when* you are in the possible world in which those uncentered propositions are true. For instance, you might know that the actual world is one with two beings in it, one who lives on the tallest mountain and one who lives on the coldest mountain, and yet not know which being *you* are.[36] And yet it seems that you might have rational credences in each of these "centered" possibilities, so the question arises of what the rational constraints on these credences are.

This question, and various answers to it, have been precisified in discussions of the "Sleeping Beauty Problem."[37] In this case, Beauty is told on Sunday that she is about to be put to sleep, and that a fair coin will be flipped (which she will not see). If the coin lands heads, she will be awakened only on Monday. If the coin lands tails, she will be awakened on Monday, then have her memory of the Monday waking erased, then be put back to sleep, and then be awakened again on Tuesday. Later, she is awakened. Question: When she is awakened, what should her credence be that the coin landed heads? Some philosophers ("Halfers") have defended the answer "one-half," whereas others ("Thirders") have defended the answer "one-third." When Beauty awakes, there are now two centerings of the tails-world that Beauty regards to be possible ("The coin landed tails and it's Monday" and "The coin landed tails and it's Tuesday"), and only one centering of the heads-world ("The coin landed heads and it's Monday") that she regards to be possible. According to the Halfer position, this is no reason for Beauty to increase her credence in the tails-world from $\frac{1}{2}$ to $\frac{2}{3}$ (and hence no reason to reduce her credence in the heads-world from $\frac{1}{2}$ to $\frac{1}{3}$). By contrast, the "Thirder" position entails that Beauty *should* increase her credence in the tails-world in response to the fact that she is now in a situation where there are two centerings that are compatible with the tails-world (and where there is still only one centering that is compatible with the heads-world).

29.5 The Requirement of Total Evidence

The Requirement of Total Evidence (RTE) is plausibly interpreted to have both a synchronic and a diachronic component, so it doesn't fit completely naturally into either of the previous two sections. The RTE enjoins us to always consider *all* of our evidence, rather than just some incomplete part of it. The synchronic constraint generated by the RTE is the constraint that the credences that a rational agent has any any time t should be the credences that are justified by his total evidence at time t. Suppose, for instance, that it's rational for you to believe at t that any particular healthy person is unlikely to die in the next year, but that healthy people *who often skydive* are likely to die in the next year. If you know that John is a healthy skydiver, then your credence that John is going to die in the next year should have the higher value justified by your total evidence which includes the fact that John skydives, rather then the lower value that would be justified by the less-than-total evidence which includes the fact that John is healthy but excludes the fact that John is a skydiver.

[36] This example comes from Lewis 1979.

[37] See, e.g., Elga 2000, Lewis 2001, Dorr 2002, Arntzenius 2002, Hitchcock 2004, and Meacham 2008.

The diachronic constraint imposed by the RTE is the constraint that your *revisions* of your credences should be the ones that are justified by a consideration of your *total* new evidence, rather than some partial component of it. This constraint is often discussed in the context of the Rule of Conditionalization; the idea is that, when you acquire some new evidence, your new credence in each A should be your old conditional credence $p(A|E)$, where E represents a *total* statement of the new evidence that you've acquired. So, for instance, when you know nothing about John and then learn that John is a healthy skydiver (and nothing more), your new credence that John will die in the next year should be your old conditional credence that John will die in the next year, conditional on his being a healthy skydiver (rather than conditional only on his being healthy, or only on his being a skydiver). But the Requirement of Total Evidence is in principle separable from the Rule of Conditionalization; even if some other update rule is correct, we still might want to insist that rational agents use a total statement of their new evidence as an input to that update rule.

Though the application of the RTE to the skydiver case is clear enough, there are some cases where its applicability is less clear. Suppose that I'm trying to assess your ability to throw a dart into the bullseye of a standard dartboard, and suppose that you hit point p inside the bullseye. It's at least somewhat natural to think that, when I update my beliefs about your dart-throwing ability, what is relevant is just the fact that your dart hit *some point or other inside the bullseye*, not that it was *point p in particular* that your dart hit. Similarly, some approaches to statistics entail that, when some experimental outcome E is observed, what is relevant is how likely it was that *E or some outcome at least as extreme as E would occur*, according to the "null hypothesis" that is being tested.[38] These approaches seem to be in some tension with the RTE, since when we know that E occurred, "E occurred" is a more complete statement of our evidence than "E or some outcome at least as extreme as E occurred."

29.6 Evidential Probability

There are two distinct philosophical theories that go by the name "evidential probability," one due primarily to Henry Kyburg (Kyburg 1961, Kyburg and Teng 2001) and one due to Timothy Williamson (Williamson 1998 and 2000).

Kyburg's theory of evidential probability was based on the idea that probabilities should be determined by relative frequencies. Evidential probability is a kind of conditional probability; the evidential probability of a sentence χ is evaluated relative to a set of sentences Γ_δ, which represents background knowledge, including knowledge of the proportions of objects satisfying various "reference class" predicates which also satisfy various "target class" predicates. The evidential probability of χ given Γ_δ, $\text{Prob}(\chi, \Gamma_\delta)$, is interval-valued, in order to accommodate cases where Γ_δ does not specify precise statistical information about the proportion of objects satisfying some reference class predicate which also satisfy the target class predicate. For example, suppose that that the proportion of red

[38] See Howson and Urbach 1993 for an overview of "frequentist" approaches to statistical inference which have this feature.

balls in an urn is known only to be between 20% and 30% inclusive, and that ball *b* is drawn from the urn. Then, Γ_δ will include sentences expressing each of these facts, and if we do not know anything *further* of relevance about *b* or the urn, then Prob(Red(*b*), Γ_δ)=[.2,.3]. (In a case where Γ_δ specifies that *precisely* 25% of the balls in the urn are red, Prob(Red(*b*), Γ_δ)=[.25,.25].) Of course, an individual object might belong to several different reference classes, each of which might contain a different proportion of individuals satisfying a particular target class predicate; for example, a individual might be both a ball in the urn and also a *plastic* ball in the urn, and the known proportion of red balls among the balls in the urn might be different from the known proportion of red balls among the plastic balls in the urn. This generates an instance of the “problem of the reference class,” which Kyburg developed a procedure for solving. In the example above, Kyburg’s theory entails that Prob(Red(*b*), Γ_δ) is equal to the known proportion of red balls among the *plastic* balls in the urn (rather than the known proportion of red balls among the balls in the urn, if this proportion is different), since the known proportion of red balls among the plastic balls is a more *specific* statistical statement. In addition to this “Specificity” principle, Kyburg also defended principles of “Richness” and “Strength,” which can be used to solve reference class problems in cases where Specificity alone doesn’t provide a unique solution.[39]

In Williamson 1998 and 2000, Timothy Williamson developed a distinct notion of evidential probability. To characterize evidential probabilities, Williamson assumes an “initial probability distribution” *P*, which is a probability function in the sense articulated in Section 3.1 above, and which measures “something like the intrinsic plausibility of hypotheses prior to investigation; this notion of intrinsic plausibility can vary in extension between contexts.”[40] On Williamson’s view, the evidential probability of *H* on total evidence *E* is $P(H|E)$; in other words, it is the conditional probability that the *P* function assigns to *H*, conditional on *E*. Evidential probabilities are distinct from credences. On Williamson’s understanding, the existence of an evidential probability for *H*, given *E*, of *x* does not entail that anyone’s credence in *H* actually is *x*. After all, *E* might not in fact be anyone’s total evidence; furthermore, present evidence might not count at all in favor of *H*, and yet everyone might still be irrationally certain of *H*. Similarly, evidential probabilities are distinct from objective physical chances or objective frequencies; a law of nature, for instance, might have an objective physical chance of 1 even though present evidence tells against its truth.

One consequence of this view is that any agent’s total evidence itself has evidential probability 1 for her; $P(E|E) = 1$ whenever it is defined. In this sense, an agent’s total evidence is certain for her. However, while the assumption that the agent updates only by the Rule of Conditionalization entails that any proposition to which an agent attaches credence 1 retains a credence of 1 forever,[41] Williamson is interested to avoid the result that a proposition that has evidential probability 1 for an agent at a time (such as the agent’s total evidence at that time) must *retain* evidential probability 1 for her permanently. One reason

[39] See Wheeler and J. Williamson 2011 for a complete discussion of Kyburg’s theory.

[40] Williamson 2000, p. 210.

[41] Suppose that $p(H) = 1$. Since $p(H|E) = \frac{p(H \wedge E)}{p(E)}$ and since $p(H) = 1$ entails that $p(H \wedge E) = p(E)$, it follows that $p(H|E) = \frac{p(E)}{p(E)} = 1$; thus, an agent who updates only by the Rule of Conditionalization will retain a credence of 1 in *H*, regardless of what evidence she acquires.

for this is that an agent might *forget* some proposition; though it was certain for her at some earlier time t_1, it is no longer certain for her at a later time t_2.

Let P be the initial probability distribution, let E_α be the the total evidence of someone in situation α, and let $P_\alpha(H)$ be the evidential probability of H for someone in α. On Williamson's view, $P_\alpha(H) = P(H|E_\alpha) = P(H \wedge E_\alpha)/P(E_\alpha)$, where $P(E_\alpha) > 0$. This allows the evidential probability of E_α to decrease from 1; in situation α, the agent's evidential probability for E_α is 1, but if the agent forgets something and her situation changes from α to $\alpha*$, then $E_\alpha*$ need not entail E_α, so it is possible that $P_{\alpha*}(E_\alpha) < 1$ even though $P_\alpha(E_\alpha) = 1$.

Here are two questions that might reasonably be asked about Williamsonian evidential probabilities. First: Do they assume the Uniqueness Thesis, according to which there is a unique rational credence to assign to a given proposition on the basis of a fixed body of total evidence? Though Williamson's use of the construction "*the* evidential probability of H on E" suggests that it might, it seems as though this assumption could be relaxed; if there is range of rationally permissible responses to a particular body of evidence, it's not obvious that there is anything in Williamson's theory that prevents the evidential probability of H on E from taking on a range of values. Second: is something like the Principal Principle (see Section 3.4) supposed to hold for evidential probabilities? In other words, suppose that X is the proposition that the *evidential probability* (rather than the *objective chance*) of A on evidence E is x; is it plausible that the agent's conditional probability for A on X and E, $p(A|X \wedge E)$, should be equal to x? Williamson doesn't address this question explicitly, but it seems quite plausible that he does understand evidential probabilities to obey this principle; after all, evidential probability is supposed to correspond to an objective notion of the probability conferred on some hypothesis by a given body of evidence, and so it seems as though a rational agent should assign a conditional credence to A, conditional on E and the information that the E confers an evidential probability of x on A, of x. Note that, unlike the case of the original Principal Principle, there can't be any such thing as inadmissible evidence here; any evidence that an agent had would be part of E.

Williamson's theory of evidential probability is very different from Kyburg's; whereas Kyburg was interested to derive evidential probabilities from statistical information alone, Williamson's theory applies to *any* body of total evidence. Both theories are also distinct from Carnap's theory of logical probabilities. Carnap's theory deserves the name "logical probability" because he thinks of probabilities as deriving from *syntactic* features of various descriptions of the domain; this is not a feature of Williamson's theory. In Carnap 1950 he assumes a domain of individuals and a number of monadic predicates that these individuals may or may not satisfy. A state description is a maximally specific description which settles whether or not each of the individuals in the domain satisfies each of the predicates under consideration. State descriptions are then grouped into equivalence classes (called "structure descriptions") under relabelings of the individuals in the domain; for example, since the state description which entails that a satisfies predicate F but that neither b nor c do can be "relabeled" to form the state description according to which b satisfies F but neither a nor c do, these two state descriptions belong to the same structure description (i.e., the structure description which says that exactly one thing is F). According to Carnap's theory, equal probability should be assigned to each structure description, and the various state descriptions that correspond to a particular structure description should be assigned an equal share of the probability assigned to that structure description. Thus, homogeneous state descriptions (such as "All of a, b, and c have F" or "None of a, b, or c has F") get higher probabilities than non-homogeneous state descriptions (such as "*a* and *b* have *F*, but *c* doesn't have *F*"), since homogeneous structure descriptions are compatible with fewer state descriptions than non-homogeneous structure descriptions are. One consequence of this, for Carnap, is that the observation of an individual satisfying *F* constitutes evidence for the homogeneous state description according to which everything is *F*. (Carnap later dropped the assumption that there is one uniquely correct inductive method, and instead proposed his "continuum of inductive methods," each member of which corresponds to learning from experience at a different "rate." See Carnap 1952 and 1963b, and Sandy Zabell's article "Symmetry Arguments in Probability" (15) in this volume (Zabell 2014).)

What unites these three distinct theories is the goal of characterizing an objective notion of probability that corresponds to how likely some hypothesis is *on a given body of evidence*, even though each theory understands the nature of that "body of evidence" differently.

29.7 Sharp and Fuzzy Credences

As discussed in Section 3.1, Probabilism entails that a rational agent's credal state is modeled by a real-valued probability function. One worry that we might have about this feature of Probabilism is that it seems psychologically unrealistic; it is hard to see what could make it the case that my credence that (say) my dog will live past 10 years old is .549123425, rather than .549123426 or .549123424. Also, it has struck some philosophers as implausible that every set of total evidence *E* rationally mandates one particular credence function; sometimes, evidence seems to be equivocal in a way that fails to rationally constrain an agent to have one particular credal state. One possible response would be to adopt a "permissive" epistemology, according to which an agent with total evidence *E* is (at least sometimes) rationally permitted to have any of multiple different credence functions.[42] A different response is to relax the requirement that a rational agent's credal state be modeled by a *single* probability function, and to allow instead that an agent's credal state be modeled by a *set* of probability functions, called a Representor.[43] If there is more than one probability function in an agent's Representor, and those functions assign different values to a proposition *A*, then an agent's credence in *A* can't be expressed as a single value; rather, the agent's credence in *A* is represented as the set of all of the values that at least one of the functions in his Representor assigns to *A*. Thus, on this view, my total evidence might permit (or even require) me to have a credence in the proposition that my dog will live past ten years old that is specified by the range [.5,.6]. In such a situation, say that my credence in *A* is "fuzzy."[44]

Despite the appeal of fuzzy credences, there are some worries that we might have about them. First, if it is psychologically unrealistic to think that I have a precise point-valued credence in the proposition that my dog will live past ten years old, it seems equally

[42] Roger White explains this view in White 2005, though he ultimately rejects it.

[43] For a more detailed summary of this sort of strategy, see Fabio Cozman's article "Imprecise Probabilities" in this volume (Cozman 2014).

[44] Various terms have been associated with this phenomenon. White 2010 refers to such credences as "mushy," Joyce 2005 as "indefinite," van Fraassen 1990 as "vague," Levi 1974 as "indeterminate," Walley 1991 as "imprecise," and Sturgeon 2008 as "thick."

psychologically unrealistic to think that I have a fuzzy credence in that proposition that is represented by precisely the interval [.5,.6], rather than the interval [.500001,.600001] or [.499999,.599999]; thus, some sort of theory of "higher-order" fuzziness seems to be required. Secondly, White 2010 argues that there is a tension between the rational permissibility of fuzzy credences and the Reflection principle (see Section 3.3); he considers a case where (on the assumption that fuzzy credences are rationally permissible) it seems as though I'm certain that my credence in *A* will rationally be fuzzy in the future, and yet where it's implausible that my credence in *A* rationally should be fuzzy now.[45] Thirdly, Elga 2010 argues that, regardless of his credence in *H*, any rational agent should be disposed to accept at least one of the following bets:

Bet A: If *H* is true, you lose $10. Otherwise you win $15.
Bet B: If *H* is true, you win $15. Otherwise you lose $10.

But, Elga argues, an agent with a fuzzy credence in *H* might be rationally permitted to reject both bets; if the range of his credence in *H* is sufficiently spread out, his credence in *H* may permit him to reject both Bet A and Bet B, and thereby forgo the guaranteed $5 that would result from accepting both bets.

29.8 Likelihood Arguments

Likelihood Arguments are arguments of a particular style that have been applied in various places in epistemology. Likelihood Arguments proceed by appealing to the Likelihood Principle, which says *E* supports *H*1 over *H*2 just in case *H*1 makes *E* likelier than *H*2 does. Suppose, for instance, that *H*1 is the proposition that the lottery is fair, *H*2 is the proposition that the lottery is heavily rigged in Mary's favor, and *E* is the proposition that Mary wins the lottery three times in a row. *H*1 seems to make *E* quite unlikely; a fair lottery is very unlikely to result in three wins in a row for anyone. But *H*2 looks to make *E* rather more likely; if the lottery is heavily rigged in Mary's favor, then it wouldn't be at all surprising for Mary to win three times. According to the Likelihood Principle, it follows that Mary's three wins are evidence for the hypothesis of rigging over the hypothesis of fairness.

The Fine-Tuning Argument is one instance of the Likelihood Argument strategy. According to the Fine-Tuning Argument, if the fundamental constants of the universe were set by chance or some other mindless natural process, it would have been very unlikely that they would all just happen to take on values that are in the narrow "life-sustaining range" and hence that can support life.[46] By contrast, if the fundamental constants were set by an Intelligent Designer who was interested in seeing to it that the universe is capable of supporting life, then it would have been far more likely that the constants would take on values in the life-sustaining range. Thus, according to the Likelihood Principle, the fact that

[45] White 2010, pp. 175–177.

[46] For a general philosophical overview of Fine-Tuning Arguments, see Leslie 1989, Holder 2004, Sober 2004, and the essays in Manson 2003. For a discussion of the physical details of Fine-Tuning Arguments, see Barrow and Tipler 1986, Collins 2003, Holder 2004 chapter 3, and Ellis section 9.1. For some objections to Fine-Tuning Arguments, see McGrew et al. 2005, Narveson 2003, and Sober 2009.

the actual constants have values in the life-sustaining range is evidence for the hypothesis of Design over the hypothesis of Chance. Likelihood Arguments have also been applied to the related question of whether the fact that the constants are life-sustaining is evidence for the Multiple Universe Hypothesis, according to which there are many non-interacting universes, each with its own set of fundamental constants. White 2003 argues that it is not; though the Multiple Universe Hypothesis does make it likelier than the Single Universe Hypothesis makes it that there would be *some* life-friendly universe, it does not make it likelier that *our* universe would be life-friendly. Bradley 2009 and Manson and Thrush 2003 develop the opposite view.

29.9 Dogmatism

In Pryor 2000, James Pryor developed an influential account of perceptual justification that he calls “dogmatism.” According to dogmatism, an agent’s perceptual experiences can give him prima facie justification to believe the basic contents of those experiences, even if he doesn’t have any “antecedent” justification to believe that his perceptual faculty is reliable, or that he isn’t a brain-in-a-vat, etc.; all that is required, according to dogmatism, is that the agent *lack* any antecedent justification to believe that such hypotheses are false. So, according to dogmatism, even if I lack justification to believe (say) that I’m not a handless brain-in-a-vat, an experience as of a hand in front of me can give me prima facie justification to believe that I have hands; and if that prima facie justification is undefeated, then my experience as of a hand can give me all-things-considered justification to believe that I have hands.

Roger White has raised the following objection to dogmatism.[47] Call the hypothesis that I am a handless brain-in-a-vat being stimulated to have hand-experiences S, and call the hypothesis that I have hands H. Suppose that, at t_1, I’m not justified in believing $\neg S$. According to dogmatism, as long as I’m also not justified at t_1 in believing S, my hand-experience E can justify me in believing H. Suppose that it does, and consider some time, t_2, after I’ve had E, at which I’m justified in believing H. Since I know that H entails $\neg S$, and I’m justified at t_2 in believing H, I should be able to competently deduce $\neg S$ from H, thereby becoming justified in believing $\neg S$ at t_2 too. But, White argues, this is quite odd. The only thing that has happened between t_1 and t_2 is that I have had experience E. But the proposition that I have E is *entailed* by S; if I am a handless brain-in-a-vat being stimulated to have hand-experiences, then I am certain to have E. And, White continues, it’s very implausible that an experience entailed by S can take me from a situation at t_1 where I am not justified in believing $\neg S$, to a situation at t_2 where I am justified in believing $\neg S$. After all, S “said” that E would occur, so E’s occurrence can’t count *against* S by justifying a belief in $\neg S$. One way of supporting this point is by appealing to the Likelihood Principle discussed in Section 7. If S entails E, then S makes E maximally likely; but if (as the Likelihood Principle entails) E is evidence for $\neg S$ over S just in case $\neg S$ makes E likelier than S does, then E can’t be evidence for $\neg S$ over S, since no hypothesis could make E any likelier than S does

[47] See White [125]. Stephen Schiffer makes a related, but distinct, objection in Schiffer 2004.

(and a fortiori ¬S can't). But if E isn't evidence for ¬S over S, then it is hard to see how I could be justified in believing ¬S at t_2, given that I wasn't justified in believing ¬S at t_1. But, according to the dogmatist, this scenario is possible.

Various responses to White's objection have been pursued. It has been generally agreed that White's objection is successful in demonstrating that dogmatism is incompatible with a probabilistic model of credences, so the responses on behalf of dogmatism have placed the blame for this incompatibility on that probabilistic model, rather than dogmatism. Pryor himself has proposed a non-probabilistic framework for credences which is based on Dempster-Shafer functions and which is designed to avoid White's objection.[48] Weatherson 2007 argues that White's objection can be avoided by a "fuzzy" model of credences of the sort discussed in Section 6.

29.10 Transmission

In various places, Crispin Wright, Martin Davies, James Pryor, and others have discussed the question of whether various arguments exhibit the phenomenon of *transmission-failure*.[49] Consider some evidence E, a hypothesis $H1$, and another hypothesis $H2$, and suppose that the relevant agent knows that $H1$ entails $H2$. Wright argues that in some cases, E provides an epistemic warrant for $H1$ and, in virtue of the fact that the agent knows that $H1$ entails $H2$, E thereby provides an epistemic warrant for $H2$. In such cases, we say that there is transmission-success. Wright [131], [133] thinks that the following case exhibits transmission-success:

Zebra
E: My experience is as of a zebra in a pen in front of me.
$H1$: There is a zebra in a pen in front of me.
$H2$: There is an animal in a pen in front of me

In Zebra [130], Wright claims, E can provide the agent with a warrant to believe $H1$, which can then "transmit" through the known entailment from $H1$ to $H2$, and thereby also provide the agent with a warrant to believe $H2$.

By contrast, Wright claims, some arguments exhibit transmission-failure. Consider:

Zebra*
E: My experience is as of a zebra in a pen in front of me.
$H1$: There is a zebra in a pen in front of me.
$H2^*$: It is not the case that there is a mule cleverly disguised to look like a zebra in a pen in front of me.

Wright argues that in Zebra*, even though E does provide a warrant for the agent to believe $H1$, and even though the agent knows that $H1$ entails $H2^*$, E does *not* provide a warrant for

[48] Pryor sets out this framework in [82], and provides some further philosophical motivation for the framework in [88]. For an excellent overview of Dempster-Shafer Theory, see Weisberg 2011 section 1.2.3 and Weisberg [120].

[49] Wright 2000, Wright 2002, Wright 2003, Wright 2004, Davies 2003, Davies 2004, Pryor 2004, Pryor 2011, Pryor forthcoming.

the agent to believe $H2^*$. The reason, Wright argues, is that in Zebra*, E provides a warrant to believe $H1$ only because the subject *already* had *independent* warrant to believe $H2^*$; thus, E cannot provide any *new* warrant to believe $H2^*$. By contrast, in Zebra, since the warrant that E provides the agent to believe $H1$ does *not* depend on her already having independent warrant to believe $H2$, the warrant that E provides for $H1$ *is* able to "transmit" through the known entailment, and thereby provide the agent with warrant to believe $H2$.

There have been a variety of attempts to translate Wright's discussion—as well as his strategy for distinguishing arguments exhibiting transmission-success from those exhibiting transmission-failure—into probabilistic terms. One such attempt is due to Okasha 2004. In cases of transmission-failure, Okasha argues, background conditions make it such that E does provide evidence for $H1$ on the assumption that $H2$ is true, so $p(H1|E \wedge H2) > p(H1|H2)$. But, since the evidence that E provides for $H1$ depends on the agent's having independent reason to believe $H2$, background conditions make it such that E does *not* provide evidence for $H1$ if we do not assume that $H2$ is true, so $p(H1|E) \leq p(H1)$. Okasha then shows that, together with the assumption that $H1$ entails $H2$, these conditions entail that $p(H2|E) \leq p(H2)$; in other words, E does not provide any evidence for $H2$.

Chandler 2010 and Moretti 2012 raise various objections to Okasha's proposal, and propose probabilistic reconstructions of their own of Wright's argument. Other proposals are developed in Kotzen 2012 and Moretti and Piazza 2012. Tucker 2010 develops a view which entails that Zebra* does exhibit transmission-success after all.

References

Appiah, A. (1985) *Assertion and Conditionals*. Cambridge: Cambridge University Press.

Armendt, B. (1993) Dutch Book, Additivity, and Utility Theory. *Philosophical Topics*. 21. pp. 1–20.

Arntzenius, F. (2002) Reflections on Sleeping Beauty. *Analysis*. 62. pp. 53–62.

Arntzenius, F. (2003) Some Problems for Conditionalization and Reflection. *Journal of Philosophy*. 100. 7. pp. 356–70.

Barrow, J. and Tipler, F. (1986) *The Anthropic Cosmological Principle*. New York, NY: Oxford University Press.

Bernstein, A. and Wattenberg, F. (1969) Non-standard measure theory. In Luxemburg, W. (ed.) *Applications of Model Theory of Algebra, Analysis, and Probability*. New York, NY: Holt, Rinehard and Winston.

Bradley, D. (2009) Multiple Universes and Observation Selection Effects. *American Philosophical Quarterly*. 46. pp. 61–72.

Bronfman, A. (ms.) A Gap in Joyce's Argument for Probabilism.

Carnap, R. (1950) *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.

Carnap, R. (1952) *The Continuum of Inductive Methods*. Chicago, IL: University of Chicago Press.

Carnap, R. (1963a) Carnap's Intellectual Autobiography. In Schilpp (ed.). *The Philosophy of Rudolf Carnap. The Library of Living Philosophers*. Vol. XI, La Salle, IL: Open Court.

Carnap, R. (1963b) Replies and Systematic Expositions. In. *The Philosophy of Rudolf Carnap. The Library of Living Philosophers*. Vol. XI, Schilpp. La Salle, IL: Open Court.

Chandler, J. (2010) The Transmission of Support: a Bayesian Re-analysis. *Synthese*. 176. pp. 333–43.

Christensen, D. Clever Bookies and Coherent Beliefs. *The Philosophical Review*. 100. pp. 229–47.
Christensen, D. (1996) Dutch-Book Arguments Depragmatized: Epistemic Consistency for Partial Believers. *Journal of Philosophy*. 93. pp. 450–79.
Christensen, D. (2004) *Putting Logic in Its Place. Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
Collins, R. (2003) Evidence for Fine-Tuning. In Manson (2003). pp. 178–99.
Cozman, F. (2016) Imprecise and Indeterminate Probabilities. In Hájek, A. and Hikheck, C. (eds.) *The Oxford handbook of Probability and Philosophy.* Oxford: Oxford University Press.
Davies, M. (2003) The Problem of Armchair Knowledge. In Nuccetelli, S. (ed.) *New Essays on Semantic Externalism and Self-Knowledge*. pp. 23–156. Cambridge, MA: MIT Press.
Davies, M. (2004) Epistemic Entitlement, Warrant Transmission and Easy Knowledge. *Proceedings of the Aristotelian Society*. Supplementary Volume 78. pp. 213–45.
de Finetti. (1937), B La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institut Henri Poincaré*. 7. pp. 1–68.
Dorr, C. (2002) Sleeping Beauty: in Defence of Elga. *Analysis*. 62. pp. 292–6.
Easwaran, K. (2011) Bayesianism I: Introduction and Arguments in Favor. *Philosophy Compass*. 6. 5. pp. 312–20.
Edwards, W., Lindman, H., and Savage, L. J. (1963) Bayesian Statistical Inference for Psychological Research. *Psychological Review*. 70. pp. 193–242.
Elga, A. (2000) Self-locating Belief and the Sleeping Beauty Problem. *Analysis*. 60. 143–7.
Elga, A. (2010) Subjective Probabilities Should be Sharp. *Philosophers' Imprint*. 10. 5.
Ellis, G. (2006) Issues in the Philosophy of Cosmology. [online] Available from http://arxiv.org/abs/astro-ph/0602280v2 [Accessed 1 Oct 2015.]
Foley, R. (1993) *Working Without a Net*. Oxford: Oxford University Press.
Frankish, K. Partial Belief and Flat-out Belief. In Huber, F. and Schmidt-Petri, C. eds. *Degrees of Belief*. pp. 75–93. Berlin: Springer.
Greaves, H. and Wallace, D. (2006) Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*. 115. pp. 607–32.
Hájek, A. (2003) What Conditional Probabilities Could Not Be. *Synthese*. 137. pp. 273–323.
Hájek, A. (2008) Arguments for—or Against—Probabilism. *The British Journal for the Philosophy of Science*. 59. 4. pp. 793–819.
Hájek, A. (2012) Is Strict Coherence Coherent?. *Dialectica*. 66. 3. pp. 411–24.
Hájek, A. (unpublished) A Puzzle about Partial Belief.
Hájek, A. (unpublished) Staying Regular.
Hall, N. (1994). Correcting the Guide to Objective Chance. *Mind*. 103. pp. 505–17.
Halpin, J. (1994) Legitimizing Chance: The Best-System Approach to Probabilistic Laws in Physical Theory. *Australasian Journal of Philosophy*. 72. pp. 317–38.
Hitchcock, C. (2004) Beauty and the Bets. *Synthese*. 139. pp. 405–20.
Holder, R. (2004) *God, the Multiverse, and Everything: Modern Cosmology and the Argument from Design*. Burlington, VT: Ashgate.
Howson, C. and Urbach, P. (1993) *Scientific Reasoning: The Bayesian Approach, 2nd Edition*. Chicago: Open Court.
Hunter, D. (1996) On the Relation Between Categorical and Probabilistic Belief. *Noûs*. 30. pp. 75–98.
Jackson, F. (1987) *Conditionals.* Oxford: Blackwell.
Jeffrey, R. (1965) *The Logic of Decision*. Chicago, IL: University of Chicago Press.

Jeffrey, R. (1970) Dracula meets Wolfman: Acceptance vs. Partial Belief. In Swain, M. (ed.) *Induction, Acceptance, and Rational Belief*. pp. 157–85. Dordrecht: D. Reidel Publishing Company.

Jeffrey, R. (1992) *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.

Jeffreys, H. (1961) *Theory of Probability.* 3rd Edition. Oxford: Clarendon Press.

Joyce, J. (1998) A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science*. 65. pp. 575–603.

Joyce, J. (1999) *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Joyce, J. (2004) Bayesianism. In Mele and Rawlings (eds) *The Oxford Handbook of Rationality*, pp. 132–55. Oxford: Oxford University Press.

Joyce, J. (2005) How Probabilities Reflect Evidence. *Philosophical Perspectives.* 19. pp. 153–78.

Joyce, J. (2009) Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Huber, F. and Schmidt-Petri, C. (eds) *Degrees of Belief*. Dordrecht: Springer.

Kaplan, M. (1996) *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.

Kemeny, J. (1955) Fair Bets and Inductive Probabilities. *Journal of Symbolic Logic.* 20. pp. 263–73.

Kolmogorov, A.(1956) *Foundations of the Theory of Probability*. 2nd Edition. New York: Chelsea Publishing Company.

Kotzen, M. (2012) Dragging and Confirming. *The Philosophical Review.* 121. pp. 55–93.

Kyburg, H. (1961) *Probability and the Logic of Rational Belief.* Middleton, CT: Wesleyan University Press.

Kyburg, H. and Teng, C. (2001) *Uncertain Inference*. Cambridge: Cambridge University Press.

Leslie, J. (1989) *Universes*. New York, NY: Routledge.

Levi, I. (1974) On Indeterminate Probabilities. *Journal of Philosophy.* 71. pp. 391–418.

Lewis, D. (1979) Attitudes De Dicto and De Se. *The Philosophical Review.* 88. pp. 513–43.

Lewis, D. (1980) A Subjectivist's Guide to Objective Chance. In Richard Jeffrey C. (ed.) *Studies in Inductive Logic and Probability*, Vol. II, pp. 263–93. Berkeley, CA: University of California Press.

Lewis, D. (1986) Introduction. In Lewis, D. *Philosophical Papers*, Vol. 2, pp. ix–xvii. Oxford: Oxford University Press.

Lewis, D. (1994) Chance and Credence: Humean Supervenience Debugged. *Mind.* 103. pp. 473–90.

Lewis, D. (1999) Why Conditionalize? In Lewis, D. *Papers in Metaphysics and Epistemology*. pp. 403–7. Cambridge: Cambridge University Press.

Lewis, D. (2001) Sleeping Beauty: Reply to Elga. *Analysis.* 61. pp. 171–76.

Loewer, B. (2004) David Lewis's Humean Theory of Objective Chance. *Philosophy of Science.* 71. pp. 1115–25.

Maher, P. (1993) *Betting on Theories*. Cambridge: Cambridge University Press.

Maher, P. (2002) Joyce's Argument for Probabilism. *Philosophy of Science.* 69. pp. 73–81.

Manson, N. (ed.) (2003) *God and Design: The Teleological Argument and Modern Science*. New York, NY: Routledge.

Manson, N. and Thrush. (2003), M Fine-Tuning, Multiple Universes, and the 'This Universe' Objection. *Pacific Philosophical Quarterly.* 84. pp. 67–83.

McGrew, T., McGrew, L. and Vestrup, E. (2001) Probabilities and the Fine-Tuning Argument: A Sceptical View. *Mind.* 110. 440. pp. 1027–37

Meacham, C. (2008) Sleeping Beauty and the Dynamics of De Se Beliefs. *Philosophical Studies.* 138. pp. 245–69.

Mellor, D. H. (1971) *The Matter of Chance*. Cambridge: Cambridge University Press.
Moretti, L. (2012) Wright, Okasha and Chandler on Transmission Failure. *Synthese*. 184. pp. 217–34.
Moretti, L. and Piazza, T. (2013) When Warrant Transmits and When It Doesn't: Towards a General Framework. *Synthèse*. 190. 13. pp. 2481–2503.
Moss. (2013), S Epistemology Formalized. *The Philosophical Review*. 122. 1. pp. 1–43.
Narveson, J. (2003) God by Design?. In Manson. pp. 88–104.
Okasha, S. (2004) Wright on the Transmission of Support: a Bayesian Analysis. *Analysis*. 64. pp. 139–46.
Perry, J. (1979) The Problem of the Essential Indexical. *Noûs*. 13. pp. 3–21.
Pryor, J. (2000) The Skeptic and the Dogmatist. *Noûs*. 34. pp. 517–49.
Pryor, J. (2004) What's Wrong with Moore's Argument? *Philosophical Issues*. 14. pp. 349–78.
Pryor, J. (2007a) What's Wrong with McKinsey-Style Reasoning? In Goldberg, S. (ed.) *Internalism and Externalism in Semantics and Episkmology* Oxford: Oxford University Press.
Pryor, J. (2007) Uncertainty and Undermining. [online] Available from: http://www.jimpryor.net/research/papers/Uncertainty.pdf. [Accessed 1 Oct 2015]
Pryor, J. (2012) Problems for credulism [online] Available from: http://www.jimpryor.net/research/papers/Credulism.pdf. [Accessed 1 Oct 2015]
Pryor, J. (2012) When Warrant Transmits. In Coliva, A. (ed.) *Wittgenstein, Epistemology and Mind: Themes from the Philosophy of Crispin Wright*. Oxford: Oxfor University Press.
Quine, W. V. (1969) Propositional Objects. In *Ontological Relativity and Other Essays*. New York, NY: Columbia University Press.
Ramsey, F. (1926) Truth and Probability. In Braithwaite, R. B. (ed.) *The Foundations of Mathematics and Other Logical Essays*, pp. 156–98. London: Kegan, Paul, Trench, Trubner & Co.; New York, NY: Harcourt, Brace and Company.
Roberts, J. (2001) Undermining Undermined: Why Humean Supervenience Never Needed to Be Debugged (Even If It's a Necessary Truth). *Philosophy of Science*, 68. 3. Supplement Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 2001) pp. S98–S108.
Savage, L. (1954) *The Foundations of Statistics*. New York, NY: John Wiley and Sons.
Schick, F. (2003) *Ambiguity and Logic*. Cambridge: Cambridge University Press.
Schiffer, S. (2004) Skepticism and the Vagaries of JustiÞed Belief. *Philosophical Studies*. 119. pp. 161–84.
Schwarz, W. (2014) Best Systems Approaches to Chance. In Hájek, A and Hikhcock C. (eds.) *The Oxford Handbook of Probability and Philosophy*. Oxford: Oxford University Press.
Seidenfeld, T. (1985) Calibration, Coherence, and Scoring Rules. *Philosophy of Science*. 52. pp. 274–94.
Shimony, A. (1955) Coherence and the Axioms of Confirmation. *Journal of Symbolic Logic*. 20. pp. 1–28.
Shimony, A. (1970) Scientific Inference. In Colodny, R. (ed.) *The Nature and Function of Scientific Theories*. Pittsburgh, PA: University of Pittsburgh Press.
Shimony, A. (1988) An Adamite Derivation of the Calculus of Probability. In Fetzer, J. H. (ed.) *Probability and Causality*. pp. 151–61. Dordrecht: Reidel.
Skyrms, B. (1980) Cousal Necessity. New Haren, CT: Yale University Press.
Skyrms, B. (1984) *Pragmatism and Empiricism*. New Haren, CT: Yale University Press.
Skyrms, B. (1995) Strict Coherence, Sigma Coherence, and the Metaphysics of Quantity. *Philosophical Studies*. 77.1. pp. 39–55.

Sober, E. (2004) The Design Argument. In Mann, W. E. (ed.) *Blackwell Guide to Philosophy of Religion*. pp. 117–47, New York, NY: Blackwell Publishers.
Sober, E. (2009) Absence Of Evidence And Evidence Of Absence: Evidential Transitivity In Connection With Fossils, Fishing, Fine-Tuning, and Firing Squads. *Philosophical Studies.* 143. pp. 63–90.
Stalnaker, R. (1970). Probability and Conditionals. *Philosophy of Science*. 37. pp. 64–80.
Stalnaker, R. (1981) Indexical Belief. *Synthese.* 49. pp. 129–51.
Stalnaker, Robert. (1984). *Inquiry*. Cambridge, MA: MIT Press.
Sturgeon, S. (2008) Reason and the Grain of Belief. *Noûs.* 42. pp. 139–65.
Talbott, W. (1991) Two Principles of Bayesian Epistemology. *Philosophical Studies.* 62. pp. 135–50.
Teller, P. (1973) Conditionalization and Observation. *Synthese.* 26. pp. 218–58.
Thau, M. (1994) Undermining and Admissibility. *Mind.* 103. pp. 491–503.
Titelbaum, M. (2016) Self-Locating Credences. In
Tucker, C. (2010) When Transmission Fails. *Philosophical Review.* 119. pp. 497–529.
van Fraassen, B. (1983) Calibration: A Frequency Justification for Personal Probability. In Cohen, R. and Laudan, L. (eds.) *Physics Philosophy and Psychoanalysis*. pp. 295–319. Dordrecht: D. Reidel.
van Fraassen, B. (1984) Belief and the Will. *Journal of Philosophy.* 81. pp. 235–56.
van Fraassen, B. (1990) Figures in a Probability Landscape. In Dunn, J. and Gupta, A. (eds.) *Truth or Consequences*. Dordrecht: Kluwer.
van Fraassen, B. (1995) Belief and the Problem of Ulysses and the Sirens. *Philosophical Studies.* 77. pp. 7–37.
Vranas, P. (1998) Who's Afraid of Undermining? Paper read at the Biennial Meeting of the Philosophy of Science Association.
Walley, P. (1991) *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
Weatherson, B. (2005) Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*. 19. 1. pp. 417–43.
Weatherson, B. (2007) The Bayesian and the Dogmatist. *Proceedings of the Aristotelian Society.* 107. pp. 169–85.
Weisberg, J. (2011) Varieties of Bayesianism. In Gabbay D., Hartmann S. and Woods, J. (eds.) *Handbook of the History of Logic*. Vol 10. Oxford: North Housand.
Weisberg, J. (2010) Dempster-Shafer Theory. [online] Available from: http://www.utm.utoronto.ca/~weisber3/unpublished/NIP\%20-\%20DST.pdf [Accessed 10 Oct 2015.]
Wheeler, G. and Williamson, J. (2011) Evidential Probability and Objective Bayesian Epistemology. In Bandyopadhyay and Forster (eds.) *Philosophy of Statistics* Vol. 7. Oxford: Elsevier.
White, R. (2003) Fine-Tuning and Multiple Universes. In Manson pp. 229–50.
White, R. (2005) Epistemic Permissiveness. *Philosophical Perspectives.* 19. pp. 445–59.
White, R. (2006) Problems for Dogmatism. *Philosophical Studies.* 131. pp. 525–57.
White, R. (2010) Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology*. Vol. 3. Oxford: Oxford University Press.
Williamson, T. (1998) Conditionalizing on Knowledge. *British Journal for the Philosophy of Science.* 49. pp. 89–121.
Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
Williamson, T. (2007) How Probable Is an Infinite Sequence of Heads? *Analysis.* 67. pp. 173–80.

Williamson, J. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.
Wright, C. (1985) Facts and Certainty. *Proceedings of the British Academy.* 71. pp. 429–72.
Wright, C. (2000) Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof. *Philosophical Issues.* 10. pp. 140–63.
Wright, C. (2002) (Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell. *Philosophy and Phenomenological Research*. 65. pp. 330–48.
Wright, C. (2003) Some Reflections on the Acquisition of Warrant by Inference. In Nuccetelli, S. (ed.) *New Essays on Semantic Externalism and Self-Knowledge*. pp. 57–77. Cambridge, MA: MIT Press.
Wright, C. (2004) Warrant for Nothing (and Foundations for Free?). *Proceedings of the Aristotelian Society*. Supplementary Vol. 78. pp. 167–212.
Zabell, S. (2014) Symmetry Arguments in Probability. In
Zynda, (2000) Representation theorems and realism about degrees of belief. *Philosophy of Science*. 67. 1. pp. 45–69.
Zynda, L. (2016) Subjective Probability. In.